Noisy Tensor Completion for Tensors with a Sparse Canonical Polyadic Factor

Swayambhoo Jain, Alexander Gutierrez, and Jarvis Haupt

arXiv:1704.02534v1 [stat.ML] 8 Apr 2017

Abstract

In this paper we study the problem of noisy tensor completion for tensors that admit a canonical polyadic or CANDECOMP/PARAFAC (CP) decomposition with one of the factors being sparse. We present general theoretical error bounds for an estimate obtained by using a complexity-regularized maximum likelihood principle and then instantiate these bounds for the case of additive white Gaussian noise. We also provide an ADMM-type algorithm for solving the complexity-regularized maximum likelihood problem and validate the theoretical finding via experiments on synthetic data set.

Index Terms

Tensor decomposition, noisy tensor completion, complexity-regularized maximum likelihood estimation, sparse CP decomposition, sparse factor models.

I. Introduction

The last decade has seen enormous progress in both the theory and practical solutions to the problem of *matrix completion*, in which the goal is to estimate missing elements of a matrix given measurements at some subset of its locations. Originally viewed from a combinatorial perspective [1], it is now usually approached from a statistical perspective in which additional structural assumptions (e.g., low-rank, sparse factors etc) not only make the problem tractable but allow for provable error bounds from noisy measurements [2]–[8]. Tensors, which we will view as multi-way arrays, naturally arise in slew of practical applications in the areas of signal processing, computer vision, neuroscience, etc. [9], [10]. Often in practice tensor data is collected in a noisy environment and suffers from missing observations. Given the success of matrix completion methods, it is no surprise that recently there has been a lot of interest in extending the successes of matrix completion to tensor completion problem [11]–[13].

In this work we consider the general problem of tensor completion. Let $\underline{X}^* \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ be the tensor we wish to estimate and suppose we collect the noisy measurements $Y_{i,j,k}$ at subset of its location $(i,j,k) \in \mathcal{S} \subset [n_1] \times [n_2] \times [n_3]$. The goal of tensor completion problem is to estimate the tensor $\underline{X}^*$ from noisy observations $\{\underline{Y}_{i,j,k}\}_{(i,j,k)\in\mathcal{S}}$. This problem is naturally ill-posed without any further assumption on the tensor we wish to estimate. We focus on structured tensors that admit *"sparse CP decomposition"* by which we mean that one of the canonical polyadic or CANDECOMP/PARAFAC (CP)-factors (defined in section I-D) is sparse. Tensors admitting

SJ and JH are with the Department of Electrical and Computer Engineering, and AG is with the School of Mathematics at University of Minnesota, Twin Cities. Author emails are:{jainx174, alexg, jdhaupt}@umn.edu. An abridged version of this paper is accepted for publication at IEEE International Symposium on Information Theory (ISIT) held in Aachen, Germany during June 25-30, 2017.

such structure arise in many applications involving electroencephalography (EEG) data, neuroimaging using functional magnetic resonance imaging (MRI), and many others [13]–[17].

A. Our Contributions

Our main contribution is encapsulated by Theorem 1 which provides general estimation error bounds for noisy tensor completion via complexity-regularized maximum likelihood estimation [18], [19] for tensors fitting our data model. This theorem can be instantiated for specific noise distributions of interest, which we do for the case when the observations are corrupted with additive white Gaussian noise. We also provide a general ADMM-type algorithm which solves an approximation to the problem of interest and then provide numerical evidence validating the statistical convergence rates predicted by Theorem 1.

B. Relation with existing works

A common theme of recent tensor completion works is modifying the tools that have been effective in tackling the matrix completion problem to apply to tensors. For example, one could apply matrix completion results to tensors directly by matricizing the tensors along various modes and minimizing the sum or weighted sum of their nuclear norms as a convex proxy for tensor rank [20]–[22]. Since the nuclear norm is computationally intractable for large scale data, matrix completion via alternating minimization was extended to tensors in [23], [24].

In contrast to these works, in this paper we consider the noisy completion of tensors that admit a CP decomposition with one of the factors being sparse. Recently, the completion of tensors with this model was exploited in the context of time series prediction of incomplete EEG data [13]. Our work is focussed on providing recovery guarantees and a general algorithmic framework and draws inspiration from recent work on noisy matrix completion under a sparse factor model [8] and extends it to tensors with a sparse CP factor.

C. Outline

After an overview of the notation used in this paper in section II we present the problem setup. In section III we present our main theorem and instantiate it for the case of Gaussian noise. In section IV we provide the algorithmic framework to solve the complexity regularized maximum likelihood estimation. Numerical experiments are provided in section VI, followed by a brief discussion and future research directions in section VII.

D. Notation

Given two continuous random variables $X \sim p(x)$ and $Y \sim q(y)$ defined on the same probability space and with p absolutely continuous with respect to q, we define the Kullback-Leibler divergence (KL-divergence) of q from p to be

$$D(p\|q) = \mathbb{E}_p\left[\log \frac{p}{q}\right].$$

If p is not absolutely continuous with respect to p, then define $D(p\|q) = \infty$. The Hellinger affinity of two distributions is similarly defined by

$$A(p,q) = \mathbb{E}_p\left[\sqrt{\frac{q}{p}}\right] = \mathbb{E}_q\left[\sqrt{\frac{p}{q}}\right].$$

We will denote vectors with lower-case letters, matrices using upper-case letters and tensors as underlined upper-case letters (e.g., $v \in \mathbb{R}^n, A \in \mathbb{R}^{m\times n}$, and $\underline{X} \in \mathbb{R}^{n_1\times n_2\times n_3}$, respectively). Furthermore, for any vector (or matrix) $v \in \mathbb{R}^n$ define $\|v\|_0 = |\{i : v_i \neq 0\}|$ to be the number of non-zero elements of v and $\|v\|_\infty := \max_i \{|v_i|\}$ to denote maximum absolute of v. Note that $\|A\|_\infty := \max_{i,j}\{|A_{i,j}|\}$ is *not* the induced norm of the matrix A.

Entry (i,j,k) of tensor $\underline{X}$ will be denoted by $X_{i,j,k}$. For a tensor $\underline{X}$ we define its Frobenius norm in analogy with the matrix case as $\|\underline{X}\|_F^2 = \sum_{i,j,k} X_{i,j,k}^2$ the squared two norm of its vectorization and its maximum absolute entry as $\|\underline{X}\|_\infty = \max_{i,j,k} |\underline{X}_{i,j,k}|$. Finally, we define the canonical polyadic or CANDECOMP/PARAFAC (CP) decomposition of a tensor $\underline{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ to be a representation

$$\underline{X} = \sum_{f=1}^{F} a_f \circ b_f \circ c_f =: [A, B, C], \tag{1}$$

where a_f, b_f, and c_f are the f^{th} columns of A, B, and C, respectively, $a_f \circ b_f \circ c_f$ denotes the tensor outer product such that $(a_f \circ b_f \circ c_f)_{i,j,k} = (i^{th}$ entry of $a_f) \times (j^{th}$ entry of $b_f) \times (k^{th}$ entry of $c_f)$, and $[A, B, C]$ is the shorthand notation of $\underline{X}$ in terms of its CP factors. The parameter F is an upper bound on the *rank* of $\underline{X}$ (we refer the reader to [10] for a comprehensive overview of tensor decompositions and their uses). For a given tensor $\underline{X}$ and CP decomposition $[A, B, C]$ define $n_{\max} = \max\{n_1, n_2, n_3, F\}$ as the maximum dimension of its CP factors and number of latent factors.

II. Problem Setup

A. Data model

Let $\underline{X}^* \in \mathcal{X} \subset \mathbb{R}^{n_1 \times n_2 \times n_3}$ be the unknown tensor whose entries we wish to estimate. We assume that X^* admits a CP decomposition such that the CP factors $A^* \in \mathbb{R}^{n_1 \times F}$, $B^* \in \mathbb{R}^{n_2 \times F}$, $C^* \in \mathbb{R}^{n_3 \times F}$ are entry-wise bounded: $\|A^*\|_\infty \leq A_{\max}$, $\|B^*\|_\infty \leq B_{\max}$, $\|C^*\|_\infty \leq C_{\max}$. Furthermore, we will assume that C^* is sparse $\|C^*\|_0 \leq k$. Then $\underline{X}^*$ can be decomposed as follow

$$\underline{X}^* = [A^*, B^*, C^*] = \sum_{f=1}^{F} a_f^* \circ b_f^* \circ c_f^*.$$

$\underline{X}$ is also entry-wise bounded, say by $\|\underline{X}^*\|_\infty \leq \frac{X_{\max}}{2}$[1]. Such tensors have a rank upper bounded by F.

B. Observation setup

We assume that we measure a noisy version of $\underline{X}^*$ at some random subset of the entries $S \subset [n_1] \times [n_2] \times [n_3]$. We generate S via an independent Bernoulli model with parameter $\gamma \in (0, 1]$ as follows: first generate $n_1 n_2 n_3$ i.i.d. Bernoulli random variables $b_{i,j,k}$ with Prob$(b_{i,j,k} = 1) = \gamma, \forall i, j, k$ and then the set S is obtained as $S = \{(i, j, k) : b_{i,j,k} = 1\}$. Conditioned on S, in the case of an additive noise model we obtain noisy observations at the locations of S as follows

$$\underline{Y}_{i,j,k} = \underline{X}^*_{i,j,k} + n_{i,j,k}, \quad \forall (i, j, k) \in S, \tag{2}$$

where $n_{i,j,k}$'s are the i.i.d noise entries.

C. Estimation procedure

Our goal here is to obtain an estimate for full true tensor $\underline{X}^*$ using the noisy sub-sampled measurement $\underline{Y}_{i,j,k}$. We pursue the complexity-regularized maximum likelihood to achieve this goal. For this we first note that the observations $\underline{Y}_{i,j,k}$ have distribution parameterized by the entries of the true tensor $\underline{X}^*$ and the overall likelihood is given by

$$p_{\underline{X}^*_S}(\underline{Y}_S) := \prod_{(i,j,k) \in S} p_{\underline{X}^*_{i,j,k}}(\underline{Y}_{i,j,k}). \tag{3}$$

[1]The factor $1/2$ is purely for the purposes of analytical tractability.

where $p_{\underline{X}^*_{i,j,k}}(\underline{Y}_{i,j,k})$ is the pdf of observation $Y_{i,j,k}$ which depends on the pdf of the noise and is parametrized by $X^*_{i,j,k}$. We use the shorthand notation $\underline{X}_S$ to denote the entries of the tensor $\underline{X}$ sampled at the indices in S.

Using prior information that C is sparse, we regularize with respect to the sparsity of C and obtain the complexity-regularized maximum likelihood estimate $\hat{\underline{X}}$ of $\underline{X}^*$ as given below

$$\hat{\underline{X}} = \underset{\underline{X}=[A,B,C]\in\mathcal{X}}{\arg\min} \left(-\log p_{\underline{X}_S}(Y_S) + \lambda \|C\|_0\right), \tag{4}$$

where $\lambda > 0$ is the regularization parameter and $\mathcal{X}$ is a class of candidate estimates. Specifically, we take $\mathcal{X}$ to be a finite class of estimates constructed as follows: first choose some $\beta \geq 1$, and set $L_{\text{lev}} = 2^{\lceil \log_2(n_{max})^\beta \rceil}$ and construct $\mathcal{A}$ to be the set of all matrices $A \in \mathbb{R}^{n_1 \times F}$ whose elements are discretized to one of L_{lev} uniformly spaced between $[-A_{\max}, A_{\max}]$, similarly construct $\mathcal{B}$ to be the set of all matrices $B \in \mathbb{R}^{n_2 \times F}$ whose elements are discretized to one of L_{lev} uniformly spaced between $[-B_{\max}, B_{\max}]$, finally $\mathcal{C}$ be the set of matrices $C \in \mathbb{R}^{n_3 \times F}$ whose elements are either zero or are discretized to one of L_{lev} uniformly spaced between $[-C_{\max}, C_{\max}]$. Then, we let

$$\mathcal{X}' = \left\{ [A,B,C] \middle| A \in \mathcal{A}, B \in \mathcal{B}, C \in \mathcal{C}, \|\underline{X}\|_\infty \leq X_{\max} \right\} \tag{5}$$

and we let $\mathcal{X}$ be any subset of $\mathcal{X}'$.

III. Main result

In this section we present the main result in which we provide an upper bound on the quality of the estimate obtained by solving (4).

Theorem 1. *Let S be sampled according to the independent Bernoulli model with parameter $\gamma = \frac{m}{n_1 n_2 n_3}$ and let Y_S be given by* (3)*. Let Q_D be any upper bound on the maximum KL divergence between $p_{\underline{X}^*_{i,j,k}}$ and $p_{\underline{X}_{i,j,k}}$ for $\underline{X} \in \mathcal{X}$*

$$Q_D \geq \max_{\underline{X}\in\mathcal{X}} \max_{i,j,k} D\left(p_{\underline{X}^*_{i,j,k}} \| p_{\underline{X}_{i,j,k}}\right)$$

where $\mathcal{X}$ is as defined in (4)*. Then for any λ satisfying*

$$\lambda \geq 4(\beta+2)\left(1+\frac{2Q}{3}\right)\log n_{max} \tag{6}$$

the regularized constrained maximum likelihood estimate $\hat{\underline{X}}$ obtained from (4) *satisfies*

$$\begin{aligned} &\frac{\mathbb{E}_{S,Y_S}\left[-2\log(A(p_{\hat{\underline{X}}}, p_{\underline{X}^*}))\right]}{n_1 n_2 n_3} \\ &\quad \leq 3 \min_{\underline{X}\in\mathcal{X}} \left\{ \frac{D(p_{\underline{X}^*} \| p_{\underline{X}})}{n_1 n_2 n_3} + \left(\lambda + \frac{8Q_D(\beta+2)\log n_{max}}{3}\right) \right. \\ &\qquad \left. \frac{(n_1+n_2)F + \|C\|_0}{m} \right\} + \frac{8Q_D \log m}{m}. \end{aligned} \tag{7}$$

Proof: The proof appears in the appendix section IX-A. ∎

The above theorem extends the main result of [8] to the tensor case. It states a general result relating the log affinity between the distributions parameterized by the estimated tensor and the ground truth tensor. Hellinger affinity is a measure of distance between two probability distributions which can be used to get bounds on the quality of the estimate. As in [8], the main utility of this theorem is that it can be instantiated for noise distributions of interest such as Gaussian, Laplace and Poisson. Note that since the estimation procedure depends only on the likelihood term, the above theorem can also be extended to non-linear observation models such as 1-bit quantized measurements [8]. We next demonstrate the utility of the above theorem to present error guarantees when the additive noise follows a Gaussian distribution.

A. Gaussian Noise Case

We examine the implications of Theorem 1 in a setting where observations are corrupted by independent additive zero-mean Gaussian noise with known variance. In this case, the observations Y_S are distributed according to a multivariate Gaussian density of dimension $|S|$ whose mean corresponds to the tensor entries at the sample locations and with covariance matrix $\sigma^2 I_{|S|}$, where $I_{|S|}$ is the identity matrix of dimension $|S|$. That is,

$$p_{\underline{X}^*_S}(\underline{Y}_S) = \frac{1}{(2\pi\sigma^2)^{|S|/2}} \exp\left(-\frac{1}{2\sigma^2}\ \|\underline{Y}_S - \underline{X}^*_S\|_F^2\right), \tag{8}$$

In order to apply Theorem 1 we choose β as:

$$\beta = \max\left\{1, 1 + \frac{\log\left(\frac{14FA_{\max}B_{\max}C_{\max}}{X_{\max}} + 1\right)}{\log(n_{\max})}\right\} \tag{9}$$

Then, we fix $\mathcal{X} = \mathcal{X}'$, and obtain an estimate according to (4) with the λ value chosen as

$$\lambda = 4\left(1 + \frac{2Q_D}{3}\right)(\beta + 2) \cdot \log(n_{\max}) \tag{10}$$

In this setting we have the following result.

Corollary 1. *Let β be as in* (9)*, let λ be as in* (10) *with $Q_D = 2X_{\max}^2/\sigma^2$, and let $\mathcal{X} = \mathcal{X}'$. The estimate $\underline{\widehat{X}}$ obtained via* (4) *satisfies*

$$\frac{\mathbb{E}_{S,Y_S}\left[\|\underline{X}^* - \underline{\widehat{X}}\|_F^2\right]}{n_1 n_2 n_3} = \mathcal{O}\left(\log(n_{\max})(\sigma^2 + X_{\max}^2)\ \left(\frac{(n_1+n_2)F + \|C^*\|_0}{m}\right)\right). \tag{11}$$

Proof: The proof appears in appendix section IX-C. ∎

Remark 1. *The quantity $(n_1 + n_2)F + \|C^*\|_0$ can be viewed as the number of degrees of freedom of the model. In this context, we note that our estimation error is proportional to the number of degrees of freedom of the model divided by m multiplied by the logarithmic factor $\log(n_{\max})$.*

Remark 2. *If we were to ignore the multilinear structure and matricize the tensor as*

$$X^*_{(3)} = (B^* \odot A^*)(C^*)^T,$$

where $\odot$ is the Khatri-Rao product (for details of matricization refer [9]) and apply the results from [8] we would obtain the bound

$$\frac{\mathbb{E}_{S,Y_S}\left[\|\underline{X}^* - \underline{\widehat{X}}\|_F^2\right]}{n_1 n_2 n_3} = \mathcal{O}\left(\log(n_{\max})(\sigma^2 + X_{\max}^2)\left(\frac{(n_1 \cdot n_2)F + \|C^*\|_0}{m}\right)\right),$$

That is, the factor of $(n_1 + n_2)F$ in Theorem 1 has become a factor of $(n_1 \cdot n_2)F$ when matricizing, a potentially massive improvement.

IV. The Algorithmic framework

In this section we propose an ADMM-type algorithm to solve the complexity regularized maximum likelihood estimate problem in (4). We note that the feasible set $\mathcal{X}$ problem in (4) is discrete which makes the algorithm design difficult. Similar to [8] we drop the discrete assumption in order to use continuous optimization techniques. This may be justified by choosing a very large value of L_{lev} and by noting that continuous optimization algorithms, when

executed on a computer, use finite precision arithmetic, and thus a discrete set of points. Hence, we consider the design of an optimization algorithm for the following problem:

$$\begin{aligned}
&\min_{\underline{X},A,B,C} -\log p_{\underline{X}_S}(\underline{Y}_S) + \lambda \left\|C\right\|_0 \\
&\text{subject to} \quad A \in \mathcal{A}, B \in \mathcal{B}, C \in \mathcal{C}, \\
&\qquad \left\|\underline{X}\right\|_\infty \leq X_{\max}, \underline{X} = \sum_{f=1}^{F} a_f \circ b_f \circ c_f, \\
&\qquad \mathcal{A} = \left\{A \in \mathbb{R}^{n_1 \times F} : \left\|A\right\|_\infty \leq A_{\max}\right\}, \\
&\qquad \mathcal{B} = \left\{B \in \mathbb{R}^{n_2 \times F} : \left\|B\right\|_\infty \leq B_{\max}\right\}, \\
&\qquad \mathcal{C} = \left\{C \in \mathbb{R}^{n_3 \times F} : \left\|C\right\|_\infty \leq C_{\max}\right\}.
\end{aligned} \tag{12}$$

We form the augmented Lagrangian for the above problem

$$\begin{aligned}
\mathcal{L}(\underline{X}, A, B, C, \lambda) = &-\log p_{\underline{X}_S}(\underline{Y}_S) + \lambda \left\|C\right\|_0 + \\
&\frac{\rho}{2}\left\|\underline{X} - \sum_{f=1}^{F} a_f \circ b_f \circ c_f\right\|_F^2 + \lambda^T \cdot \text{vec}\left(\underline{X} - [A, B, C]\right) \\
&+ I_{\mathcal{X}}(\underline{X}) + I_{\mathcal{A}}(A) + I_{\mathcal{B}}(B) + I_{\mathcal{C}}(C),
\end{aligned}$$

where λ is Lagrangian vector of size $n_1 n_2 n_3$ for the tensor equality constraint and $I_{\mathcal{X}}(\underline{X}), I_{\mathcal{A}}(A), I_{\mathcal{B}}(B), I_{\mathcal{C}}(C)$ are indicator functions of the sets $\|\underline{X}\|_\infty \leq X_{\max}$, $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ respectively[2]. Starting from the augmented Lagrangian we extend the ADMM-type algorithm proposed in [8] to the tensor case as shown in Algorithm 1.

V. The algorithm

Algorithm 1 ADMM-type algorithm for noisy tensor completion

Inputs: $\Delta_1^{\text{stop}}, \Delta_2^{\text{stop}}, \eta, \rho^{(0)}$

Initialize: $\underline{X}^{(0)}$, $A^{(0)}$,$B^{(0)}$,$C^{(0)}$, $\lambda^{(0)}$

while $\Delta_1 > \Delta_1^{\text{stop}}, \Delta_2 > \Delta_2^{\text{stop}}, t \leq t_{max}$ **do**

S1: $\underline{X}^{(t+1)} = \arg\min_{\underline{X}} \mathcal{L}(\underline{X}, A^{(t)}, B^{(t)}, C^{(t)}, \lambda^{(t)})$

S2: $A^{(t+1)} = \arg\min_{A} \mathcal{L}(\underline{X}^{(t+1)}, A, B^{(t)}, C^{(t)}, \lambda^{(t)})$

S3: $B^{(t+1)} = \arg\min_{B} \mathcal{L}(\underline{X}^{(t+1)}, A^{(t+1)}, B, C^{(t)}, \lambda^{(t)})$

S4: $C^{(t+1)} = \arg\min_{C} \mathcal{L}(\underline{X}^{(t+1)}, A^{(t+1)}, B^{(t+1)}, C, \lambda^{(t)})$

S5: $\lambda^{(t+1)} = \lambda^{(t)} + \rho^{(0)} \text{vec}\left(\underline{X}^{(t+1)} - [A^{(t+1)}, B^{(t+1)}, C^{(t+1)}]\right)$

Set $\Delta_1 = \left\|\underline{X}^{(t+1)} - [A^{(t+1)}, B^{(t+1)}, C^{(t+1)}]\right\|_F$

Set $\Delta_2 = \rho^{(k)} \left\|[A^{(t)}, B^{(t)}, C^{(t)}] - [A^{(t+1)}, B^{(t+1)}, C^{(t+1)}]\right\|_F$

$$\rho^{(k+1)} = \begin{cases} \eta\rho^{(k)}, \text{ if} \Delta_1 \geq 10\Delta_2 \\ \rho^{(k)}/\eta, \text{ if} \Delta_2 \geq 10\Delta_1 \\ \rho^{(k)}, \text{ otherwise} \end{cases}$$

end while

Output: $A = A^{(t)}, B = B^{(t)}, C = C^{(t)}$

[2]The convex indicator of set U is defined as $I_U(x) =$ if $x \in U$ and $I_U(x) = \infty$if $x \notin U$. Note that function $I_U(x)$ is convex function if U is convex set.

The $\underline{X}$ update in Algorithm 1 is separable across components and so it reduces to $n_1 n_2 n_3$ scalar problems. Furthermore, the scalar problem is closed-form for $(i,j,k) \notin S$ and is a proximal-type step for $(i,j,k) \in S$. This is a particularly attractive feature because many common noise densities (e.g., Gaussian, Laplace) have closed-form proximal updates [8]. The A and B updates can be converted to a constrained least squares problem and can be solved via projected gradient descent. We solve the C update via iterative hard thresholding. Although the convergence of this algorithm to a stationary point remains an open question and a subject of future work, we have not encountered problems with this in our simulations.

VI. Numerical Experiments

In this section we include simulations which corroborate our theorem. For each experiment we construct the true data tensor $\underline{X} = [A^*, B^*, C^*]$ by individually constructing the CP factors A^*, B^*, C^* (as described below), where the magnitudes of entries of the true factors A^*, B^*, and C^* are bounded in magnitude by $A^*_{\max}, B^*_{\max}$, and $C^*_{\max}$ respectively. For the purposes of these experiments we fix $n_1 = 30, n_2 = 30, n_3 = 50$ and $A^*_{\max} = 1, B^*_{\max} = 1, C^*_{\max} = 10$.

For a given F the true CP factors were generated as random matrices of dimensions $n_1 \times F$, $n_2 \times F$, $n_3 \times F$ with standard Gaussian $\mathcal{N}(0,1)$ entries. We then projected the entries of the A and B matrices so that $\|A^*\|_\infty \leq A^*_{\max}$ and $\|B^*\|_\infty \leq B^*_{\max}$. For the C^* matrix we first project C^* entry-wise to the interval $[-C_{\max}, C_{\max}]$ and then pick k entries uniformly at random and zero out all other entries so that we get the desired sparsity $\|C^*\|_0 = k$. From these tensors the tensor $\underline{X}^*$ was calculated as $\underline{X}^* = [A^*, B^*, C^*]$ as in (1).

We then take measurements at a subset of entries following a Bernoulli sampling model with sampling rate $\gamma \in (0,1]$ and corrupt our measurements with additive white Gaussian noise of variance $\sigma = 0.25$ to obtain the final noisy measurements. The noisy measurements were then used to calculate the estimate by solving (an approximation to) the complexity regularized problem in (12) using algorithm 1. Note that for Gaussian noise the negative log-likelihood in problem (12) reduces to a squared error loss over the sampled entries. Since in practice the parameters $A_{\max}$, $B_{\max}$, $C_{\max}, X_{\max}$ are not known *a priori* we will assume we have an upper bound for them and in our experiments set them as $A_{\max} = 2A^*_{\max}, B_{\max} = 2B^*_{\max}, C_{\max} = 2C^*_{\max}, X_{\max} = 2\|\underline{X}^*\|_\infty$. Further, we also assume that F is known *a priori*.

In figure 1 we show how the log per entry squared error $\log\left(\frac{\|\hat{\underline{X}} - \underline{X}^*\|_F^2}{n_1 n_2 n_3}\right)$ decays as a function of log sampling rate $\log(\gamma)$ for $F = 5, 15$ in the paper and a fixed sparsity level $\|C\|_0 = 0.2 n_3 F$. The plot is obtained after averaging over 10 trials to average out random Bernoulli sampling at given sampling rate γ and noise. Each plot corresponds to a single chosen value of λ, selected as the value that gives a representative error curve (e.g., one giving lowest overall curve, over the range of parameters we considered). Our theoretical results predict that the error decay should be inversely proportional to the sampling rate $\gamma = \frac{m}{n_1 n_2 n_3}$ when viewed on a log-log scale, this corresponds to the slope of -1. The curve of $F = 5$ and $F = 15$ are shown in blue solid line and red dotted line. For both the cases the slope of curves is similar and it is approximately -1. Therefore these experimental results validate both the theoretical error bound in corollary 1 and the performance of our proposed algorithm.

VII. Conclusion and Future Directions

In this work we extend the statistical theory of complexity-penalized maximum likelihood estimation developed in [8], [18], [19] to noisy tensor completion for tensors admitting CP decomposition with a sparse factor. In particular, we provide theoretical guarantees on the performance of sparsity-regularized maximum likelihood estimation under a Bernoulli sampling assumption and general i.i.d. noise. We then instantiate the general result for the specific case of

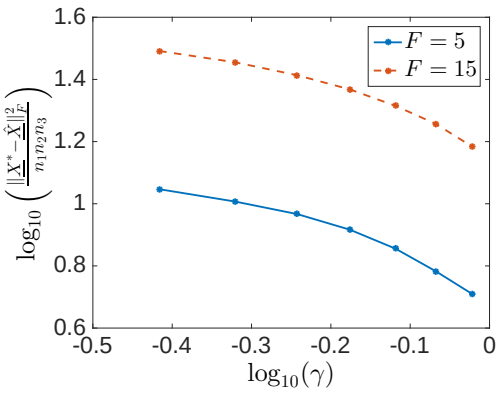

Fig. 1. Plot for log per-entry approximation error $\log\left(\frac{\|\underline{\hat{X}}-\underline{X}^*\|_F^2}{n_1n_2n_3}\right)$ vs the log sampling rate: $\log(\gamma)$ for the two ranks $F = 5, 15$. The slope at the higher sampling rates is approximately -1 (the rate predicted by our theory) in both cases.

additive white Gaussian noise. We also provided an ADMM-based algorithmic framework to solve the complexity-penalized maximum likelihood estimation problem and provide numerical experiments to validate the theoretical bounds on synthetic data.

Obtaining error bounds for other noise distributions and non-linear observation setting such 1-bit quantized observations is an interesting possible research direction. Extending the main result to approximately sparse CP factor or to tensors with multiple sparse CP factor are also important directions for future research.

VIII. Acknowledgements

We thank Professor Nicholas Sidiropoulos for his insightful guidance and discussions on tensors which helped in completion of this work. Swayambhoo Jain and Jarvis Haupt were supported by the DARPA Young Faculty Award, Grant N66001-14-1-4047. Alexander Gutierrez was supported by the NSF Graduate Research Fellowship Program under Grant No. 00039202.

IX. Appendix

A. Proof of Main Theorem

The proof of our main result is an application of the following general lemma.

Lemma 1. *Let $\underline{X}^* \in \mathbb{R}^{n_1\times n_2\times n_3}$ and let $\mathcal{X}$ be a finite collection of candidate reconstructions with assigned weights $\mathrm{pen}(\underline{X}) \geq 1$ satisfying the Kraft-McMillan inequality over $\mathcal{X}$.*

$$\sum_{\underline{X}\in\mathcal{X}} 2^{-\mathrm{pen}(\underline{X})} \leq 1. \tag{13}$$

Fix an integer $k \leq m \leq n_1n_2n_3$ and let $\gamma = \frac{m}{n_1n_2n_3}$ and generate $n_1n_2n_3$ i.i.d. Bernoulli(γ) random variables $S_{i,j,k}$ so that entry $(i,j,k) \in S$ if $S_{i,j,k} = 1$ and $(i,j,k) \notin S$ otherwise. Conditioned on S we obtain independent measurements $Y_S \sim p_{X_S^*} = \prod_{(i,j,k)\in S} p_{X^*_{i,j,k}}$. Then if Q_D is an upper bound for the maximum KL-divergence

$$Q_D \geq \max_{\underline{X}\in\mathcal{X}} \max_{(i,j,k)} D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}),$$

it follows that for any

$$\xi \geq (1 + \frac{2Q_D}{3}) \cdot 2\log 2 \tag{14}$$

the compexity-penalized maximum likelihood estimator

$$\hat{\underline{X}}^{\xi}(S, \underline{Y}_S) = \arg\min_{\underline{X} \in \mathcal{X}} \left\{ -\log p_{\underline{X}_S}(\underline{Y}_S) + \xi \operatorname{pen}(\underline{X}) \right\}$$

satisfies the error bound

$$\frac{\mathbb{E}_{S,\underline{Y}_S}\left[-2\log(A(p_{\hat{X}^*}, p_{X^*}))\right]}{n_1 n_2 n_3} \leq \frac{8 Q_D \log m}{m} + 3 \min_{\underline{X} \in \mathcal{X}} \left\{ \frac{D(p_{\underline{X}^*} \| p_{\underline{X}})}{n_1 n_2 n_3} + \left(\xi + \frac{4 Q_D \log 2}{3} \right) \frac{\operatorname{pen}(\underline{X})}{m} \right\}.$$

Proof: The proof appears in Appendix section IX-B. ■

For using the result in Lemma 1 we need to define penalties $\operatorname{pen}(\underline{X}) \geq 1$ on candidate reconstructions $\underline{X}$ of $\underline{X}^*$, so that for every subset $\mathcal{X}$ of the set $\mathcal{X}'$ specified in the conditions of Theorem 1 the summability condition $\sum_{\underline{X} \in \mathcal{X}} 2^{-\operatorname{pen}(\underline{X})} \leq 1$ holds. To this end, we will use the fact that for any $\mathcal{X} \subseteq \mathcal{X}'$ we always have $\sum_{\underline{X} \in \mathcal{X}} 2^{-\operatorname{pen}(\underline{X})} \leq \sum_{\underline{X} \in \mathcal{X}'} 2^{-\operatorname{pen}(\underline{X})}$; thus, it suffices for us to show that for the specific set $\mathcal{X}'$ described in (5), the penalty satisfies the Kraft-McMillan inequality:

$$\sum_{\underline{X} \in \mathcal{X}'} 2^{-\operatorname{pen}(\underline{X})} \leq 1. \tag{15}$$

The Kraft-Mcmillan Inequality is automatically satisfied if we set the $\operatorname{pen}(\underline{X})$ to be the code length of some uniquely decodable binary code for the elements $\underline{X} \in \mathcal{X}'$ [25].

We utilize a common encoding strategy for encoding the elements of $\mathcal{A}$ and $\mathcal{B}$. We encode each entry of the matrices using $\log_2(L_{\text{lev}})$ bits in this manner the total number of bits needed to code any elements in $\mathcal{A}$ and $\mathcal{B}$ is $n_1 F \log_2(L_{\text{lev}})$ and $n_2 F \log_2(L_{\text{lev}})$ respectively. Since the elements of set $\mathcal{C}$ are sparse we follow a two step procedure: first we encode the location of the non-zero elements using $\log_2 L_{\text{loc}}$ bits where $L_{\text{loc}} = 2^{\lceil \log_2(n_3 F) \rceil}$ and then we encode the entry using $\log_2(L_{\text{lev}})$ bits. Now, we let $\mathcal{X}''$ be the set of all such $\underline{X}$ with CPD factors $A \in \mathcal{A}$,$B \in \mathcal{B}$, $C \in \mathcal{C}$, and let the code for each $\underline{X}$ be the concatenation of the (fixed-length) code for A followed by (fixed-length) code for B followed by the (variable-length) code for C. It follows that we may assign penalties $\operatorname{pen}(\underline{X})$ to all $\underline{X} \in \mathcal{X}''$ whose lengths satisfy

$$\operatorname{pen}(\underline{X}) = (n_1 + n_2) F \log_2 L_{\text{lev}} + \|C\|_0 \log_2(L_{\text{loc}} L_{\text{lev}}).$$

By construction such a code is uniquely decodable, since by the Kraft McMillan inequality we have $\sum_{X \in \mathcal{X}''} 2^{-\operatorname{pen}(\underline{X})} \leq 1$. Further, since $\mathcal{X}' \subset \mathcal{X}''$ this also satisfies the inequality $\sum_{\underline{X} \in \mathcal{X}} 2^{-\operatorname{pen}(\underline{X})} \leq 1$ in (13) in Lemma 1 is satisfied for $\mathcal{X}'$ sa defined in statement of the Theorem 1. Now for any set $\underline{X} \subseteq \mathcal{X}'$ and using coding strategy described above, the condition (13) in Lemma (1) is satisfied. So for randomly subsampled and noisy observations $Y_{\mathcal{S}}$ our estimates take the form

$$\begin{aligned} \widehat{\underline{X}}^{\xi} &= \arg\min_{\underline{X} = [A,B,C] \in \mathcal{X}} \{ -\log p_{X_{\mathcal{S}}}(Y_{\mathcal{S}}) + \xi \operatorname{pen}(\underline{X}) \} \\ &= \arg\min_{\underline{X} = [A,B,C] \in \mathcal{X}} \{ -\log p_{X_{\mathcal{S}}}(Y_{\mathcal{S}}) + \xi \log_2(L_{\text{loc}} L_{\text{lev}}) \|C\|_0 \} \end{aligned}$$

Further, when ξ satisfies (14), we have

$$\begin{aligned}\frac{\mathbb{E}_{S,\underline{Y}_S}\left[-2\log(A(p_{\hat{\underline{X}}^*},p_{\underline{X}^*}))\right]}{n_1n_2n_3} &\leq \frac{8Q_D\log m}{m}+\\ 3\min_{\underline{X}\in\mathcal{X}}&\left\{\frac{D(p_{\underline{X}^*}\|p_{\underline{X}})}{n_1n_2n_3}+\left(\xi+\frac{4Q_D\log 2}{3}\right)\right.\\ &\left.\cdot\frac{(n_1+n_2)F\log_2 L_{\rm lev}+\|C\|_0\log_2(L_{\rm loc}L_{\rm lev})}{m}\right\}\\ &\leq \frac{8Q_D\log m}{m}+\\ 3\min_{\underline{X}\in\mathcal{X}}&\left\{\frac{D(p_{\underline{X}^*}\|p_{\underline{X}})}{n_1n_2n_3}+\left(\xi+\frac{4Q_D\log 2}{3}\right)\right.\\ &\left.\cdot\log_2(L_{\rm loc}L_{\rm lev})\frac{(n_1+n_2)F+\|C\|_0}{m}\right\}.\end{aligned}$$

Finally, we let $\lambda=\xi\cdot\log_2(L_{\rm loc}L_{\rm lev})$ and using the relation that

$$\log_2 L_{\rm loc}L_{\rm lev}\leq 2\cdot(\beta+2)\cdot\log(n_{\max}) \tag{16}$$

which follows by our selection of $L_{\rm lev}$ and $L_{\rm loc}$ and the fact that $F,n_3\leq n_{\max}$ and $n_{\max}\geq 4$. Using the condition (16) and (14) in Lemma 1 it follows that for

$$\lambda\geq 4(\beta+2)\left(1+\frac{2Q_D}{3}\right)\log(n_{\max})$$

the estimate

$$\hat{\underline{X}}^{\lambda}=\underset{\underline{X}=[A,B,C]\in\mathcal{X}}{\arg\min}\left(-\log p_{\underline{X}_S}(Y_S)+\lambda\|C\|_0\right), \tag{17}$$

satisfies the bound (7) in Thereom 1.

B. Proof of Lemma

The main requirement for the proof of this lemma is to show that our random Bernoulli measurement model is "good" in the sense that it will allow us to apply some known concentration results. Let Q_D be an upper bound on the KL-divergence of $p_{X_{i,j}}$ from $p_{\underline{X}^*_{i,j}}$ over all elements $X\in\mathcal{X}$:

$$Q_D\geq\max_{\underline{X}\in\mathcal{X}}\max_{i,j}D(p_{\underline{X}^*_{i,j,k}}\|p_{\underline{X}_{i,j,k}}).$$

Similarly, let Q_A be an upper bound on negative two times the log of the Hellinger affinities between the same:

$$Q_A\geq\max_{\underline{X}\in\mathcal{X}}\max_{i,j}-2\log\left(A(p_{\underline{X}^*_{i,j,k}}\|p_{\underline{X}_{i,j,k}})\right).$$

Let $m\leq n_1n_2n_3$ be the expected total number of measurements and $\gamma=m/(n_1n_2n_3)$ to be the ratio of measured entries to total entries. Given any $\delta\in(0,1)$ define the "good" set $\mathcal{G}_{\gamma,\delta}$ as the subset of all possible sampling sets that satisfy a desired property:

$$\begin{aligned}\mathcal{G}_{\gamma,\delta}:=\Bigg\{S\subseteq[n_1]\times[n_2]\times[n_3]:&\left(\bigcap_{\underline{X}\in\mathcal{X}}D(p_{\underline{X}^*_S}\|p_{\underline{X}_S})\leq\frac{3\gamma}{2}D(p_{\underline{X}^*}\|p_{\underline{X}})+(4/3)Q_D[\log(1/\delta)+\mathrm{pen}(\underline{X})\log 2]\right)\\ &\cap\left(\bigcap_{\underline{X}\in\mathcal{X}}(-2\log A(p_{\underline{X}^*_S},p_{\underline{X}_S}))\geq\frac{\gamma}{2}(-2\log A(p_{\underline{X}^*},p_{\underline{X}}))-(4/3)Q_A[\log(1/\delta)+\mathrm{pen}(\underline{X})\log 2]\right)\Bigg\}\end{aligned}$$

We show that an Erdós-Renyi model with parameter γ will be "good" with high probability in the following lemma.

Lemma 2. *Let $\mathcal{X}$ be a finite collection of countable estimates $\underline{X}$ for $\underline{X^*}$ with penalties $\mathrm{pen}(\underline{X})$ satifying the Kraft inequality* (13)*. Then for any fixed $\gamma, \delta \in (0,1)$ let S be a random subset of $[n_1] \times [n_2] \times [n_3]$ be a random subset generated according the Erdós-Renyi model.Then $\mathbb{P}[S \notin \mathcal{G}_{\gamma,\delta}) \leq 2\delta$.*

Proof: Note that $\mathcal{G}_{\gamma,\delta}$ is defined in terms of an intersection of two events, define them to be

$$\mathcal{E}_D := \left\{ \bigcap_{\underline{X} \in \mathcal{X}} D(p_{\underline{X^*}_S} \| p_{\underline{X}_S}) \leq \frac{3\gamma}{2} D(p_{\underline{X^*}} \| p_{\underline{X}}) + (4/3) Q_D [\log(1/\delta) + \mathrm{pen}(\underline{X}) \log 2] \right\}$$

and

$$\mathcal{E}_A := \left\{ \bigcap_{\underline{X} \in \mathcal{X}} (-2 \log A(p_{\underline{X^*}_S}, p_{\underline{X}_S})) \geq \frac{\gamma}{2} (-2 \log A(p_{\underline{X^*}}, p_{\underline{X}})) - (4/3) Q_A [\log(1/\delta) + \mathrm{pen}(\underline{X}) \log 2] \right\}.$$

We apply the union bound to find that

$$\mathbb{P}\left[S \notin \mathcal{G}_{\gamma,\delta}\right] \leq \mathbb{P}\left[\mathcal{E}_u^C\right] + \mathbb{P}\left[\mathcal{E}_\ell^C\right],$$

and will prove the theorem by showing that each of the two probabilities on the right-hand side are less than δ, starting with $\mathbb{P}[\mathcal{E}_u^C]$.

Since the observations are conditionally independent given S, we know that for fixed $\underline{X} \in \mathcal{X}$,

$$D(p_{\underline{X}^*_S} \| p_{\underline{X}_S}) = \sum_{(i,j,k) \in S} D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) = \sum_{i,j,k} S_{i,j,k} D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}),$$

where $S_{i,j,k} \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(\gamma)$. We will show that random sums of this form are concentrated around its mean using the Craig-Bernstein inequality .

The version ofthe Craig-Bernstein inequality that we will use states: let $U_{i,j,k}$ be random variables such that we have the uniform bound $|U_{i,j,k} - \mathbb{E}[U_{i,j,k}]| \leq \beta$ for all i, j, k. Let $\tau > 0$ and ϵ be such that $0 < \epsilon\beta/3 < 1$. Then

$$\mathbb{P}\left[\sum_{i,j,k} (U_{i,j,k} - \mathbb{E}[U_{i,j,k}]) \geq \frac{\tau}{\epsilon} + \epsilon \frac{\sum_{i,j,k} \mathrm{var}(U_{i,jk})}{2(1 - \epsilon\beta/3)} \right] \leq e^{-\tau}.$$

To apply the Craig-Bernstein inequality to our problem we first fix $\underline{X} \in \mathcal{X}$ and define $U_{i,j,k} = S_{i,j,k} D(p_{\underline{X}^*_{i,j,k}} \| p_{X_{i,j,k}})$. Note that $U_{i,j,k} \leq Q_D \Rightarrow |U_{i,j,k} - \mathbb{E}[U_{i,j,k}]| \leq Q_D$. We also bound the variance via

$$\begin{aligned} \mathrm{var}(U_{i,j,k}) &= \gamma(1-\gamma) \left(D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \right)^2 \\ &\leq \gamma \left(D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \right)^2. \end{aligned}$$

Then let $\epsilon = \frac{3}{4Q_D}$ and $\beta = Q_D$ in (IX-B) to get that

$$\mathbb{P}\left[\sum_{i,j,k} (S_{i,j,k} - \gamma) D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \geq \frac{4Q_D\tau}{3} + \frac{\sum_{i,j,k} \gamma \cdot \left(D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \right)^2}{2Q_D} \right] \leq e^{-\tau}.$$

Now use the fact that $D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \leq Q_D$ by definition to cancel out the square term to get:

$$\mathbb{P}\left[\sum_{i,j,k} (S_{i,j,k} - \gamma) D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \geq \frac{4Q_D\tau}{3} + \frac{3\gamma}{2} \sum_{i,j,k} \cdot D(p_{X^*_{i,j,k}} \| p_{X_{i,j,k}}) \right] \leq e^{-\tau}.$$

Finally, we define $\delta = e^{-\tau}$, and simplify to arrive at

$$\mathbb{P}\left[D(p_{\underline{X}^*_S} \| p_{\underline{X}_S}) \geq \frac{4Q_D \log(1/\delta)}{3} + \frac{3\gamma}{2} D(p_{\underline{X}^*} \| p_{\underline{X}}) \right] \leq \delta, \tag{18}$$

for any δ.

To get a uniform bound over all $\underline{X} \in \mathcal{X}$ define $\delta_{\underline{X}} := \delta 2^{-\operatorname{pen}(\underline{X})}$ and use the bound in (18) with $\delta_{\underline{X}}$ and apply the union bound over the class $\mathcal{X}$ to find that

$$\mathbb{P}\left[\bigcup_{\underline{X}\in\mathcal{X}} D(p_{\underline{X}_S^*}\|p_{\underline{X}_S}) \geq \frac{3\gamma}{2} D(p_{\underline{X}^*}\|p_{\underline{X}}) + \frac{4Q_D}{3}\left[\log(1/\delta) + \operatorname{pen}(\underline{X})\cdot\log 2\right]\right] \leq \delta. \tag{19}$$

An similar argument (applying Craig-Bernstein and a union bound) can be applied to $\mathcal{E}_A$ to obtain

$$\mathbb{P}\left[\bigcup_{\underline{X}\in\mathcal{X}} \left(-2\log A(p_{\underline{X}_S^*}, p_{\underline{X}_S})\right) \leq \frac{\gamma}{2}(-2\log A(p_{\underline{X}^*}, p_{\underline{X}})) - (4Q_A/3)[\log(1/\delta) + \operatorname{pen}(\underline{X})\cdot\log 2]\right] \leq \delta \tag{20}$$

This completes the proof of lemma 2. ∎

Given lemma 2, the rest of the proof of lemma 1 is a straightforward extension of the already-published proof of lemma A.1 in [8].

C. Proof of Corollary 1

We first establish a general error bound, which we then specialize to the case stated in the corollary. Note that for $\underline{X}^*$ as specified and any $\underline{X} \in \mathcal{X}$, using the model (8) we have

$$D(p_{X^*_{i,j,k}}\|p_{X_{i,j,k}}) = \frac{(X^*_{i,j,k} - X_{i,j,k})^2}{2\sigma^2}$$

for any fixed $(i,j,k) \in S$. It follows that $D(p_{\underline{X}^*}\|p_{\underline{X}}) = \|\underline{X}^* - \underline{X}\|_F^2/2\sigma^2$. Further. as the amplitudes of entries of $\underline{X}^*$ and all $\underline{X} \in \mathcal{X}$ upper bounded by $X_{\max}$, it is easy to see that we may choose $Q_D = 2X_{\max}^2/\sigma^2$. Also, for any $\underline{X} \in \mathcal{X}$ and any fixed $(i,j,k) \in S$ it is easy to show that in this case

$$-2\log A(p_{X_{i,j,k}}, p_{X^*_{i,j,k}}) = \frac{(X^*_{i,j,k} - X_{i,j,k})^2}{4\sigma^2},$$

so that $-2\log A(p_{\underline{X}}, p_{\underline{X}^*}) = \|\underline{X}^* - \underline{X}\|_F^2/4\sigma^2$. It follows that

$$\mathbb{E}_{S,\underline{Y}_S}\left[-2\log A(p_{\underline{\widehat{X}}}, p_{\underline{X}^*})\right] = \frac{\mathbb{E}_{S,\underline{Y}_S}\left[\|\underline{X}^* - \underline{\widehat{X}}\|_F^2\right]}{4\sigma^2}.$$

Now for using Theorem 1, we first substitute the value of $Q_D = 2X_{\max}^2/\sigma^2$ to obtain the following condition on λ

$$\lambda \geq 4\cdot\left(1 + \frac{4X_{\max}^2}{3\sigma^2}\right)\cdot(\beta+2)\cdot\log(n_{\max}).$$

Above condition implies that the specific choice of λ given (10) is a valid choice to use if we want to invoke Theorem 1. So fixing λ as given (10) and using Theorem 1, the sparsity penalized ML estimate satisfies the per-element mean-square error bound

$$\frac{\mathbb{E}_{S,\underline{Y}_S}\left[\|\underline{X}^* - \underline{\widehat{X}}\|_F^2\right]}{n_1n_2n_3} \leq \frac{64X_{\max}^2\log m}{m} + 6\cdot\min_{\underline{X}\in\mathcal{X}}\left\{\frac{\|\underline{X}^* - \underline{X}\|_F^2}{n_1n_2n_3} + \left(2\sigma^2\lambda + \frac{24X_{\max}^2(\beta+2)\log(n_{\max})}{3}\right)\left(\frac{(n_1+n_2)F + \|C\|_0}{m}\right)\right\}.$$

Notice that the above inequality is sort of an oracle type inequality because it implies that for any $\underline{X} \in \mathcal{X}$ we have

$$\frac{\mathbb{E}_{S,\underline{Y}_S}\left[\|\underline{X}^* - \underline{\widehat{X}}\|_F^2\right]}{n_1n_2n_3} \leq \frac{64X_{\max}^2\log m}{m} + 6\cdot\left\{\frac{\|\underline{X}^* - \underline{X}\|_F^2}{n_1n_2n_3} + \left(2\sigma^2\lambda + \frac{24X_{\max}^2(\beta+2)\log(n_{\max})}{3}\right)\left(\frac{(n_1+n_2)F + \|C\|_0}{m}\right)\right\}.$$

We use this inequality for a specific candidate reconstruction of the form $\underline{X}_Q^* = [A_Q^*, B_Q^*, C_Q^*]$ where the entries of s A_Q^* are the closest discretized surrogates of the entries of A^*, B_Q^* are the closest discretized surrogates of the

entries of B^*, and C_Q^* are the closest discretized surrogates of the non-zeros entries of C^* (and zero otherwise). For proceeding further we need to bound $\|\underline{X}_Q^* - \underline{X}^*\|_{\max}$. For this purpose we consider matricization of tensor across the third dimension as follows

$$\|\underline{X}_Q^* - \underline{X}^*\|_{\max} = \left\| \left(B_Q^* \odot A_Q^*\right) \left(C_Q^*\right)^T - \left(B^* \odot A^*\right) \left(C^*\right)^T \right\|_{\max}$$

Next we write $A_Q^* = A^* + \Delta_A$, $B_Q^* = B^* + \Delta_B$ and $C_Q^* = C^* + \Delta_C$ with straight forward matrix multiplication we can obtain that

$$\begin{aligned}\left(B_Q^* \odot A_Q^*\right) \left(C_Q^*\right)^T =& \left(B^* \odot A^*\right) \left(C^*\right)^T + \left(\Delta_A \odot B^* + A^* \odot \Delta_B + \Delta_A \odot \Delta_B\right) \left(C^*\right)^T \\ &+ \left(A^* \odot B^* + \Delta_A \odot B^* + A^* \odot \Delta_B + \Delta_A \odot \Delta_B\right) \Delta_C^T \end{aligned} \tag{21}$$

Using this identity it follows

$$\|\underline{X}_Q^* - \underline{X}^*\|_{\max} = \left\| \left(\Delta_A \odot B^* + A^* \odot \Delta_B + \Delta_A \odot \Delta_B\right) \left(C^*\right)^T + \left(A^* \odot B^* + \Delta_A \odot B^* + A^* \odot \Delta_B + \Delta_A \odot \Delta_B\right) \Delta_C^T \right\|_{\max}$$

Now using the facts that $\|A \odot B\|_{\max} = \|A\|_{\max}\|B\|_{\max}$, $\|AB\|_{\max} \leq F\|A\|\|B\|_{\max}$ and triangle inequality for the $\|\cdot\|_{\max}$ norm it is easy to show that

$$\|\underline{X}_Q^* - \underline{X}^*\|_{\max} \leq F\left[(\|\Delta_A\|_{\max} + \|A\|_{\max})(\|\Delta_B\|_{\max} + \|B\|_{\max})(\|\Delta_C\|_{\max} + \|C\|_{\max}) - \|A\|_{\max}\|B\|_{\max}\|C\|_{\max}\right]$$

Further, using the fact that $\|\Delta_A\|_{\max} \leq \frac{A_{\max}}{L_{\text{lev}}-1}$, $\|\Delta_B\|_{\max} \leq \frac{B_{\max}}{L_{\text{lev}}-1}$, and $\|\Delta_C\|_{\max} \leq \frac{C_{\max}}{L_{\text{lev}}-1}$, we have

$$\begin{aligned}&\|\underline{X}_Q^* - \underline{X}^*\|_{\max} \\ &\leq F\left[\left(\frac{A_{\max}}{L_{\text{lev}}-1} + A_{\max}\right)\left(\frac{B_{\max}}{L_{\text{lev}}-1} + \|B\|_{\max}\right)\left(\frac{C_{\max}}{L_{\text{lev}}-1} + C_{\max}\right) - A_{\max}B_{\max}C_{\max}\right] \\ &\leq FA_{\max}B_{\max}C_{\max}\left[\left(1 + \frac{1}{L_{\text{lev}}-1}\right)^3 - 1\right] \\ &\leq \frac{FA_{\max}B_{\max}C_{\max}}{L_{\text{lev}}-1}\left[3 + \frac{3}{L_{\text{lev}}-1} + \frac{1}{(L_{\text{lev}}-1)^2}\right] \\ &\leq \frac{7FA_{\max}B_{\max}C_{\max}}{L_{\text{lev}}-1},\end{aligned}$$

where in the second last step we have used $L_{\text{lev}} \geq 2$. Now, it is straight-forward to show that our choice of β in (9) implies $L_{\text{lev}} \geq 14FA_{\max}B_{\max}C_{\max}/X_{\max} + 1$, so each entry of $\|\underline{X}_Q^* - \underline{X}^*\|_{\max} \leq X_{\max}/2$. This further implies that for the candidate estimate $\underline{X}_Q^*$ we have $\|\underline{X}_Q^*\|_{\max} \leq X_{\max}$, i.e., $\underline{X}_Q^* \in \mathcal{X}$. Moreover, we

$$\frac{\|\underline{X}^* - \underline{X}_Q^*\|_F^2}{n_1n_2n_3} \leq \left(\frac{7FA_{\max}B_{\max}C_{\max}}{L_{\text{lev}}-1}\right)^2 \leq \frac{X_{\max}^2}{m}, \tag{22}$$

where the last inequality follows from the fact that our specific choice of β in (9) also implies $L_{\text{lev}} \geq 7F\sqrt{m}A_{\max}B_{\max}C_{\max}/X_{\max}$.

Finally, we evaluate the oracle inequality for (21) for $\underline{X}_Q^*$ and using the fact that $\|C_Q^*\|_0 = \|C^*\|_0$ and using the value of λ specified in the corollary we have

$$\frac{\mathbb{E}_{S,\underline{Y}_S}\left[\|\underline{X}^* - \widehat{\underline{X}}\|_F^2\right]}{n_1n_2n_3} \leq \frac{70X_{\max}^2 \log m}{m} + 24(\sigma^2 + 2X_{\max}^2)(\beta+2)\log(n_{\max})\left(\frac{(n_1+n_2)F + \|C^*\|_0}{m}\right).$$

REFERENCES

[1] Leslie Hogben, "Graph theoretic methods for matrix completion problems," *Linear Algebra and Its Applications*, vol. 328, no. 1-3, pp. 161–202, 2001.

[2] Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh, "Matrix completion from noisy entries," in *Advances in Neural Information Processing Systems*, 2009, pp. 952–960.

[3] Emmanuel J Candes and Yaniv Plan, "Matrix completion with noise," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 925–936, 2010.

[4] Emmanuel J Candès and Terence Tao, "The power of convex relaxation: Near-optimal matrix completion," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.

[5] Benjamin Recht, "A simpler approach to matrix completion," *The Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, 2011.

[6] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 2013, pp. 665–674.

[7] Akshay Soni, Swayambhoo Jain, Jarvis Haupt, and Stefano Gonella, "Error bounds for maximum likelihood matrix completion under sparse factor models," in *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*. IEEE, 2014, pp. 399–403.

[8] Akshay Soni, Swayambhoo Jain, Jarvis Haupt, and Stefano Gonella, "Noisy matrix completion under sparse factor models," *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3636–3661, June 2016.

[9] Tamara G Kolda and Brett W Bader, "Tensor decompositions and applications," *SIAM review*, vol. 51, no. 3, pp. 455–500, 2009.

[10] Nicholas D Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E Papalexakis, and Christos Faloutsos, "Tensor decomposition for signal processing and machine learning," *arXiv preprint arXiv:1607.01668*, 2016.

[11] Liang Xiong, Xi Chen, Tzu-Kuo Huang, Jeff G Schneider, and Jaime G Carbonell, "Temporal collaborative filtering with Bayesian probabilistic tensor factorization.," in *SDM*. SIAM, 2010, vol. 10, pp. 211–222.

[12] Kejun Huang, Nicholas D Sidiropoulos, and Athanasios P Liavas, "A flexible and efficient algorithmic framework for constrained matrix and tensor factorization," *IEEE Transactions on Signal Processing*, vol. 64, no. 19, pp. 5052–5065, 2016.

[13] Weiwei Shi, Yongxin Zhu, S Yu Philip, Mengyun Liu, Guoxing Wang, Zhiliang Qian, and Yong Lian, "Incomplete electrocardiogram time series prediction," in *Biomedical Circuits and Systems Conference (BioCAS), 2016 IEEE*. IEEE, 2016, pp. 200–203.

[14] Genevera Allen, "Sparse higher-order principal components analysis.," in *AISTATS*, 2012, vol. 15.

[15] Roland Ruiters and Reinhard Klein, "BTF compression via sparse tensor decomposition," in *Computer Graphics Forum*. Wiley Online Library, 2009, vol. 28, pp. 1181–1188.

[16] Evangelos E Papalexakis, Nicholas D Sidiropoulos, and Rasmus Bro, "From k-means to higher-way co-clustering: Multilinear decomposition with sparse latent factors," *IEEE transactions on signal processing*, vol. 61, no. 2, pp. 493–506, 2013.

[17] Yanwei Pang, Zhao Ma, Jing Pan, and Yuan Yuan, "Robust sparse tensor decomposition by probabilistic latent semantic analysis," in *Image and Graphics (ICIG), 2011 Sixth International Conference on*. IEEE, 2011, pp. 893–896.

[18] Eric D Kolaczyk and Robert D Nowak, "Multiscale likelihood analysis and complexity penalized estimation," *Annals of statistics*, pp. 500–527, 2004.

[19] Jonathan Q Li and Andrew R Barron, "Mixture density estimation," in *Advances in Neural Information Processing Systems 12*. Citeseer, 1999.

[20] Ming Yuan and Cun-Hui Zhang, "On tensor completion via nuclear norm minimization," *Foundations of Computational Mathematics*, vol. 16, no. 4, pp. 1031–1068, 2016.

[21] Ji Liu, Przemyslaw Musialski, Peter Wonka, and Jieping Ye, "Tensor completion for estimating missing values in visual data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 208–220, 2013.

[22] Bo Huang, Cun Mu, Donald Goldfarb, and John Wright, "Provable low-rank tensor recovery," *Optimization-Online*, vol. 4252, pp. 2, 2014.

[23] Yangyang Xu, Ruru Hao, Wotao Yin, and Zhixun Su, "Parallel matrix factorization for low-rank tensor completion," *arXiv preprint arXiv:1312.1254*, 2013.

[24] Prateek Jain and Sewoong Oh, "Provable tensor factorization with missing data," in *Advances in Neural Information Processing Systems*, 2014, pp. 1431–1439.

[25] Thomas M Cover and Joy A Thomas, *Elements of information theory*, John Wiley & Sons, 2012.